

REMARKS

In the non-final Office Action, the Examiner rejected claims 1, 3-11, 13, 15, 16, and 19-33 under 35 U.S.C. § 103(a) as unpatentable over Dinker et al. (U.S. Patent No. 7,206,836) in view of Brin et al. ("The Anatomy of a Large-Scale Hypertextual Web search Engine") and Rao et al. (U.S. Patent No. 5,689,706). Applicants traverse the rejection. Claims 1, 3-11, 13, 15, 16, and 19-33 remain pending.

In paragraph 4 of the Office Action, the Examiner rejected claims 1, 3-11, 13, 15, 16, and 19-33 under 35 U.S.C. § 103(a) as allegedly unpatentable over Dinker et al. in view of Brin et al. and Rao et al. Applicants traverse the rejection.

Independent claim 1, for example, is directed to a file system that comprises a plurality of servers configured to store file data as chunks; and a master connected to the servers and configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks, store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks, where the master is configured to communicate with the servers upon startup of the master to identify the chunks stored by the servers, and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data.

Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, do not disclose or suggest the combination of features recited in claim 1. For example, Dinker et al., Brin et al., and Rao et al. do not disclose or suggest a master, connected

to a plurality of servers that are configured to store file data as chunks, that is configured to, among other things, store an operation log that includes a record of changes to at least one of namespace data or mapping data, where the namespace data includes file identifiers for files for which the file data is stored as chunks, and the mapping data maps the file identifiers to the chunks to which the file identifiers correspond.

The Examiner admitted that Dinker et al. and Brin et al. do not disclose or suggest an operation log, but alleged that Rao et al. discloses a master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and cited column 9, lines 8-25, of Rao et al. for support (Office Action, page 4). Applicants submit that the disclosure of Rao et al. does not support the Examiner's allegation.

At column 9, lines 8-25, Rao et al. discloses:

Operation is as follows: when application process 503 is initialized, it obtains a file identifier from pipe process 711 which specifies pipe 710. When execution of application program 509 results in performance of a file operation, the function for the operation in lib.3d 507 both causes kernel server 305(a) to perform the function on the file system provided by kernel server 305(a) and also sends a message via pipe 710 to backup system 513. When the message arrives in backup system 513, it is received by backup log process 716, which logs the message in a log file 703(b) in the file system provided by kernel server 305(b). Whenever log file 703(b) has messages in it, the messages are read in the order in which they arrived by syscall engine process 715. In the preferred embodiment, backend map 517 belongs to syscall engine process 715. As syscall engine process 715 reads the messages, it causes kernel server 305(b) to perform the file operations required by the messages and itself maintains backend map 517 as required by the messages.

Rao et al. discloses a backup file system that includes a primary system 511 upon which application process 503 runs, and a backup system 513 upon which backup copies of the files modified by application process 503 are maintained (see col. 6, lines 5-18). In the section identified above, Rao et al. discloses that when application program 509 performs a file operation, a message is sent to backup system 513 that logs the message in a log file, reads the

message, and performs the file operations required by the message. Rao et al. does not disclose or remotely suggest that primary system 511, or any other element, is a master connected to a plurality of servers that are configured to store file data as chunks. Because Rao et al. does not disclose or remotely suggest a master that is connected to a plurality of servers that are configured to store file data as chunks, Rao et al. cannot disclose or suggest a master connected to a plurality of servers and configured to, for example, store an operation log that includes a record of changes to at least one of namespace data or mapping data, as recited in claim 1.

The Examiner alleged that it would have been obvious to use the master of Dinker et al. and Brin et al. to save the log disclosed by Rao et al. "in order to provide accurate information pertaining to the location of files by tracking location changes" (Office Action, page 4). Applicants submit that the Examiner's motivation statement lacks merit. Contrary to the Examiner's allegation, Rao et al. does not disclose or suggest storing a log that assists in tracking location changes of files. Instead, Rao et al. discloses a log file that includes messages pertaining to file operations performed by a primary system 511 (col. 9, lines 8-25). Thus, incorporating the log file of Rao et al. into a combined system of Dinker et al. and Brin et al. would not provide accurate information pertaining to the location of files by tracking location changes, as alleged by the Examiner. Therefore, the Examiner's motivation statement lacks merit and the Examiner failed to establish a prima facie case of obviousness with regard to claim 1.

Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, also do not disclose or suggest a master connected to a plurality of servers and configured to communicate with the servers upon startup of the master to identify the chunks

stored by the servers, as further recited in claim 1.

The Examiner alleged that Dinker et al. discloses these features and cited column 6, lines 8-67, of Dinker et al. for support (Office Action, page 3). Applicants submit that the disclosure of Dinker et al. provides no support for the Examiner's allegation.

At column 6, lines 8-67, Dinker et al. discloses a replication topology manager 160 that is configured to maintain the distribution of data defined by a replication topology (defined as a static definition of how data should be replicated within a cluster (col. 4, lines 20-22)). Even assuming, for the sake of argument, that replication topology manager 160 can be equated to a master and nodes 101 can be equated to servers (points that Applicants do not concede), nowhere does Dinker et al. disclose or remotely suggest that replication topology manager 160 communicates with nodes 101 upon startup of replication topology manager 160 to identify the chunks stored by nodes 101. Thus, Dinker et al. does not disclose or suggest a master connected to a plurality of servers and configured to communicate with the servers upon startup of the master to identify the chunks stored by the servers, as recited in claim 1. Brin et al. and Rao et al. also do not disclose or suggest these features.

Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, also do not disclose or suggest a master connected to a plurality of servers and configured to record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as further recited in claim 1.

The Examiner alleged that Dinker et al. discloses these features and cited column 6, lines 8-67, of Dinker et al. for support (Office Action, page 3). Applicants submit that the disclosure of Dinker et al. provides no support for the Examiner's allegation.

At column 6, lines 8-67, Dinker et al. discloses a replication topology manager 160 that is configured to maintain the distribution of data defined by a replication topology (defined as a static definition of how data should be replicated within a cluster (col. 4, lines 20-22)). Even assuming, for the sake of argument, that replication topology manager 160 can be equated to a master (a point that Applicants do not concede), nowhere does Dinker et al. disclose or remotely suggest that replication topology manager 160 maintains the distribution of data in a non-persistent manner. Thus, Dinker et al. does not disclose or suggest a master connected to a plurality of servers and configured to record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as recited in claim 1. Brin et al. and Rao et al. also do not disclose or suggest these features.

For at least these reasons, Applicants submit that claim 1 is patentable over Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination. Claims 3-10 and 19-26 depend from claim 1 and are, therefore, patentable over Dinker et al., Brin et al., and Rao et al. for at least the reasons given with regard to claim 1. Claims 3-10 and 19-26 are also patentable over Dinker et al., Brin et al., and Rao et al. for reasons of their own.

For example, claim 22 recites that a chunk handle, which uniquely identifies one of the chunks, encodes a timestamp. Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, do not disclose or suggest this feature recited in claim 22.

The Examiner alleged that Brin et al. discloses this feature and cited Section 4.2.2 of Brin et al. for support (Office Action, page 8). Applicants submit that the disclosure of Brin et al. does not support the Examiner's allegation.

In Section 4.2.2, Brin et al. discloses a repository that stores a compressed version of the

full HTML of every web page prefixed by docID, length, and URL. Brin et al. does not disclose or suggest file data that is stored as chunks by a plurality of servers. Therefore, Brin et al. cannot disclose or suggest a chunk handle, which uniquely identifies one of the chunks, that encodes a timestamp, as recited in claim 22.

The Examiner alleged that the wordID disclosed by Brin et al. corresponds to a chunk handle (Office Action, page 8). Even assuming, for the sake of argument, that the wordID can be equated to a chunk handle (a point that Applicants do not concede), Brin et al. does not disclose or remotely suggest that the wordID encodes a timestamp, as recited in claim 22.

For at least these additional reasons, Applicants submit that claim 22 is patentable over Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination.

Claim 23 recites that the master is configured to update the location data by periodically instructing the servers to provide information regarding the chunks stored by the servers. Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, do not disclose or suggest this feature recited in claim 23.

The Examiner alleged that Dinker et al. discloses this feature and cited column 6, lines 8-67, of Dinker et al. for support (Office Action, page 8). Applicants submit that the disclosure of Dinker et al. does not support the Examiner's allegation.

At column 6, lines 8-67, Dinker et al. discloses that replication topology manager 160 responds to the addition, departure, and failure of nodes in a cluster by performing one or more copy operations involving data in the data store so that the actual replication of the data within the cluster closely follows that defined by the replication topology. Even assuming, for the sake of argument, that the nodes can be equated to servers and the replication topology manager can

be equated to a master (points that Applicants do not concede), nowhere in this section, or elsewhere, does Dinker et al. disclose or suggest that the replication topology manager periodically instructs the nodes to provide information regarding the data stored by the nodes. Thus, Dinker et al. cannot disclose or suggest a master that is configured to update the location data by periodically instructing the servers to provide information regarding the chunks stored by the servers, as recited in claim 23.

For at least these additional reasons, Applicants submit that claim 23 is patentable over Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination.

Claim 24 recites that the operation log includes a logical timeline that defines an order for concurrent operations. Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, do not disclose or suggest this feature recited in claim 24.

The Examiner alleged that Rao et al. discloses this feature and cited column 9, lines 8-25, of Rao et al. for support (Office Action, page 8). Applicants submit that the disclosure of Rao et al. does not support the Examiner's allegation.

At column 9, lines 8-25, Rao et al. discloses that when application program 509 (in primary system 511) performs a file operation, a message is sent to backup system 513 that logs the message in a log file, reads the message, and performs the file operations required by the message. Nowhere in this section, or elsewhere, does Rao et al. disclose or suggest concurrent operations, let alone an operation log that includes a logical timeline that defines an order for concurrent operations, as recited in claim 24.

For at least these additional reasons, Applicants submit that claim 24 is patentable over Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination.

Claim 25 recites that the master is configured to determine when a size of the operation log exceeds a threshold, and create a checkpoint of the operation log when the size of the operation log exceeds the threshold. Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, do not disclose or suggest these features recited in claim 25.

The Examiner alleged that Rao et al. discloses these features and cited column 11, lines 11-34, of Rao et al. for support (Office Action, page 8). Applicants submit that the disclosure of Rao et al. does not support the Examiner's allegation.

At column 11, lines 11-34, Rao et al. discloses:

Syscall engine 715 fails: Monitor 719 detects the failure and restarts syscall engine 715. With checkpoint and recovery functions provided by libft, the newly restarted syscall engine 715 is able to recover to its previously-checkpointed status from an external file. No other processes are aware of this failure/recovery.

Backend log process 716 fails: Monitor 719 detects the failure and restarts backend log process 716. Again, process 716 restores its status from a checkpoint file. Monitor 719 further informs monitor 717 that backend log process 716 has been restarted, and monitor 717 in turn informs pipe process 711. Process 711 then connects pipe 710 to the new backend log process 716. The next write of each application fails and lib.3d gets the new connection from pipe process 711

Front-end log process 705 fails: Front-end log process 705 exists only during periods of failure of system 513. When monitor 717 detects the failure of front-end log process 705, it informs pipe process 711. which then restarts front-end log process 705 and reconnects pipe 708 to it. The next write of application program 509 fails and the message sending function in lib.3d gets the file descriptor for the new pipe 708 from pipe process 711.

In this section, Rao et al. discloses that when syscall engine 715 fails, syscall engine 715 is restarted from its previously-checkpointed status from an external file; when backend log process 716 fails, backend log process 716 is restarted from a checkpoint file; and when front-end log process 705 fails, front-end log process 705 is restarted. Nowhere in this section, or elsewhere, does Rao et al. disclose or remotely suggest a master that is configured to determine when a size of an operation log exceeds a threshold, as recited in claim 25.

For at least these additional reasons, Applicants submit that claim 25 is patentable over Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination.

Independent claims 13, 15, 16, and 30 recite features similar to, yet possibly different in scope from, features recited in claim 1. Claims 13, 15, 16, and 30 are, therefore, patentable over Dinker et al., Brin et al., and Rao et al., whether taken alone or in any reasonable combination, for at least reasons similar to reasons given with regard to claim 1. Claims 27-29 depend from claim 13, and claims 31-33 depend from claim 30. Claims 27-29 and 31-33 are, therefore, patentable over Dinker et al., Brin et al., and Rao et al. for at least the reasons given with regard to claims 13 and 30, respectively.

In view of the foregoing remarks, Applicants respectfully request the Examiner's reconsideration of the application and the timely allowance of pending claims 1, 3-11, 13, 15, 16, and 19-33.

As Applicants' remarks with respect to the Examiner's rejections are sufficient to overcome these rejections, Applicants' silence as to certain assertions by the Examiner in the Office Action or certain requirements that may be applicable to such rejections (e.g., whether a reference constitutes prior art, motivation to combine references, assertions regarding dependent claims, etc.) is not a concession by Applicants that such assertions are accurate or such requirements have been met, and Applicants reserve the right to analyze and dispute these assertions/requirements in the future.

If the Examiner does not believe that all pending claims are now in condition for allowance, the Examiner is urged to contact the undersigned to expedite prosecution of this application.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, LLP

By: /Paul A. Harrity, Reg No 39574/
Paul A. Harrity
Reg. No. 39,574

Date: January 31, 2008

11350 Random Hills Road
Suite 600
Fairfax, Virginia 22030
(571) 432-0800

Customer Number 44989